Investigating Strategies for Managing Civil Violence using the MANA Agent Based Distillation

Sau Yee Yiu[1], Andrew Gill[2] and Peng Shi[1]

Land Operations Division[1]
Defence Systems Analysis Division[2]
Defence Science and Technology Organisation (DSTO)
PO Box 1500, Edinburgh SA 5111, Australia

ABSTRACT: A remarkably simple model of civil violence has recently been built and studied by Joshua Epstein at the Center of Social and Economic Dynamics in the US using a very simple cellular automata (CA) simulation. However, the model and its analysis were based on the assumption that the entities have purely random movement, which limits the degree of realism of the model. DSTO, through LOD and DSAD, has access to a more sophisticated CA model known as MANA, developed by the Defence Technology Agency in NZ. Recently, DSTO has extended the original US model and analysis to incorporate various movement strategies of the entities by using the MANA simulation. This paper describes the model and the analysis of the data, which included graphical, statistical and game theory techniques, and which provide some initial thoughts on the effectiveness of various strategies for managing civil violence. These results may also have applicability in other Operations Other Than War (OOTW) scenarios, including peace keeping and counter-terrorism. Various possible extensions to the extended MANA model and analysis are also covered.

1. Introduction

1.1 Epstein Civil Violence Model

In this highly idealized model [1], a central authority seeks to suppress a decentralized rebellion. The model contains "*Quiets*" (members of the general population), "*Actives*" (those *quiets* who have become actively rebellious), and "*Cops*" (forces of the central authority who seek out and arrest actively rebellious agents). All entities possess local vision (modelled by a finite radius) and move randomly over a 2D lattice (representing some region). Such models are known as cellular automata models.

Therefore, over time, members of the general population can transition to various states (see Figure 1 **Q** = *Quiets*, **A** = *Actives*, **J** = *Jail*) and it was the central authority's aim to maximise the number of *quiets*, while the rebel's aim was to maximise the number of *actives*.

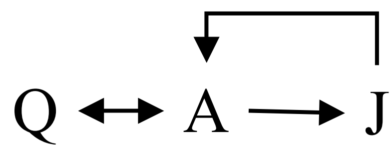

Figure 1. Epstein Model Transitions

The rules that govern the transitions are as follows: the *quiets (actives)* will become *active (quiet)* if their (estimated) probability of being arrested is less (greater) than a certain threshold. This estimate is assumed to increase with the ratio of *cops* to *actives* within the prospective rebel's vision, and was given by:

$$\mathbf{P = 1\text{-}exp[-k(C/A)]} \qquad \mathbf{(1)}$$

where **exp** stands for exponential function, **k** is a constant to be set, **C** is the number of *Cops* and **A** is the number of *Actives*. The constant k is set to ensure a plausible estimate (say P=0.9) when C=1 and A=1.

For a fixed number of *cops*, the agent's estimated arrest probability falls the more *actives* there are, and this simple idea played an important role in the analysis.

The *cops* (which never defect to the revolution in this model) had one simple rule, which was to inspect all sites within its vision and randomly arrest an *active*. An arrested *active* was then released after a finite duration and was assumed to be active.

The question Epstein explored was whether this highly idealized model was *sufficient to generate* recognizable macroscopic revolutionary dynamics of fundamental interest, to which the answer was yes.

1.2 Motivation for Current Study

The dynamics of the Epstein model (transitions between *quiets* and *actives* and *jailed*) are essentially governed by the fluctuating spatial densities of the *quiets*, *actives* and *cops* as they randomly move about the region. For example, pockets of relatively low *cops* densities are 'ripe' for rebellious activities and these can grow unchecked, as observed by the former China president, Mao Tse Tung ("a single spark can cause a prairie fire"), and which is why freedom of assembly is often the first casualty of repressive regimes as suggested by Epstein [1].

Given the importance of relative spatial densities, it is therefore conceivable that each side might improve their chances of success if they adopt some movement strategy that is not random. It is also somewhat artificial to assume the *cops* do not react (probably by chasing) to detected *actives* and vice versa.

DSTO has been using a more sophisticated cellular automata model, known as MANA [2], which importantly includes non-random movement strategies. By using the MANA simulation [3], DSTO, through LOD [4] and DSAD [5] has extended the original Epstein Civil Violence model to incorporate various movement strategies (random and non-random) of the *active* and *cops* entities within the model.

This paper describes the extended Epstein model, comments on the effectiveness of various movement strategies for managing civil violence, and, finally makes suggestions on possible further extensions to the model.

A broader motivation for this study is the observation that DSTO does not possess a robust set of modelling and analysis tools for OOTW type scenarios (including peace keeping, criminal and terrorist networks) although defence planning for these operations are becoming more frequently required. This study may shed some light on whether cellular automata models are an appropriate enabler for this type of analysis.

2. MANA Civil Violence Model

2.1 Modelling State Transitions

There are two differences in the state transition diagram (see Figure 2) from the Epstein model. The first is that there is no direct transition from *active* back to *quiet*. This represents a somewhat more zealous *active* entity than in the Epstein model and will provide more of a test of the *cops'* strategy. The second difference is that *jailed actives* are assumed to be *quiet* on release, which is meant to represent some form of rehabilitation.

$$\begin{tikzcd} Q \arrow[r] & A \arrow[r] & J \arrow[ll, bend right] \end{tikzcd}$$

Figure 2. MANA Model Transitions

The MANA modelling of the *active-to-jailed* state transition is essentially the same as in the original Epstein model, in that a single (randomly selected) *active* within the *cop's* vision is arrested. The *jailed-to-quiet* transition is also easily modelled in MANA by assigning a duration to remain in the *jailed* state and assigning the *quiet* state as the fallback state at the end of the duration.

However, MANA is incapable of modelling the *quiet-to-active* transition as governed by the arrest probability defined in equation (**1**). The scheme used in MANA to approximate this situation is governed by the entity interaction probabilities.

To simulate a higher probability of a *quiet* becoming *active* when in a high *active-to-cops* density region, we assign a probability of a successful *cops-to-active* interaction to imply a temporary subduing of the *active* (with a subdued *active* unable to interact), and a probability of a successful *active-to-quiet* interaction to imply a transition to *active* of the *quiet*. Thus, in a high *active-to-cops* density region, there will be relatively more unsubdued *actives* that will then have a relatively higher probability of transitioning a *quiet* to the *active* state.

2.2 Movement Strategies

Movement of the *quiet* entities is assumed to be random (as in the Epstein model) for two reasons. First, we assume the *quiets* have no affinity to either the *active* or *cop* entities. Second, we wish to concentrate on the dynamics between the *active* and *cops* movement strategies.

Also, movement for *cops* and *actives* are simulated in MANA by a set of weightings, which describe an entity's propensity to move toward or away from other *cops*, *actives* and/or *quiets*. To allow representation of a strategy by a scalar parameterisation, we define:

$$\lambda = -\frac{\mathbf{W}_{\mathbf{Cops}}}{\mathbf{W}_{\mathbf{Quiets}}} \qquad \mu = \frac{\mathbf{W}_{\mathbf{Actives}}}{\mathbf{W}_{\mathbf{Quiets}}}$$

as the movement strategies for the *actives* and *cops*, respectively. In this study, we assume that $\mathbf{W}_{\mathbf{Cops}}$ is negative and $\mathbf{W}_{\mathbf{Actives}}$ and $\mathbf{W}_{\mathbf{Quiets}}$ are both positive. That is, *actives* are repelled from *cops*, *cops* are attracted to *actives*, and both may have some attraction to the *quiets*.

Furthermore, we assume that

$$-\mathbf{W}_{\mathbf{Cops}} + \mathbf{W}_{\mathbf{Quiets}} = \mathbf{1}$$

for the *actives*, and that

$$\mathbf{W}_{\mathbf{Actives}} + \mathbf{W}_{\mathbf{Quiets}} = \mathbf{1}$$

for the *cops*, which models a trade-off between avoiding (or chasing) the opposition and inciting (or protecting) the general population.

For example, a large value of λ and small value of μ would represent a situation where *actives* are very cautious of the *cops* but the *cops* are more interested in protecting the general population.

2.3 Baseline Scenario

Figure 3 on the next page shows a screen shot of the baseline scenario. The scenario begins with twelve *cops*, and a population of two hundred, of which twenty are actively rebellious and the remainder are *quiets*. The map area represents a grid of 200 squares by 200 squares.

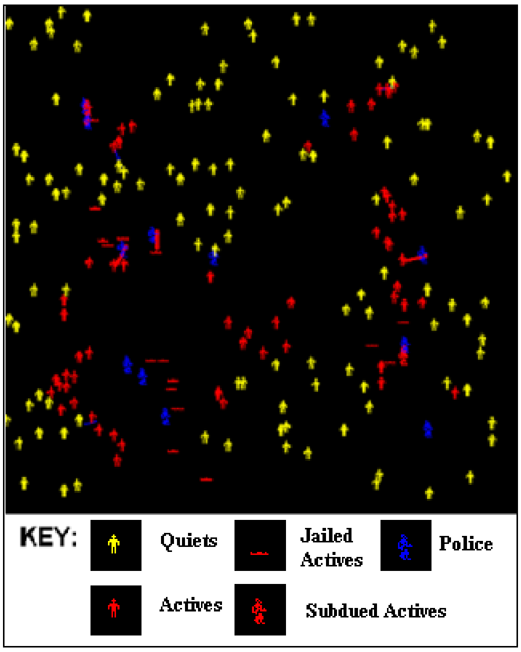

Figure 3. Baseline Mana Civil Violence Scenario

Table 1 lists a summary of each squad properties. *Jailed actives* are detained for a fixed term of 50 time-steps during which they play no part in the scenario.

	Cops	*Actives*	*Quiets*
Number of Agents	12	20	180
Sensor Range	50	50	10
Interaction Range	10	10	0
Probability Subduing Actives	50	0	0
Probability Converting Quiets	0	10	0

Table 1. Summary of Squad Properties

3. Analysis of Data

In the first part of the analysis, we will investigate the effectiveness of various movement strategies for both the *Cops* and *Active* as defined by the parameters λ and μ using the baseline MANA scenario. From the results of the resulting payoff matrix we can deduce the preferred strategies for both the *Cops* and *Actives*.

Using these preferred strategies, we then study the impacts of the more resource demanding options of increased jail times and number of *cops* on our civil violence model.

Each simulation is run to 2000 time-steps to allow the system to reach equilibrium and the mean percentage of *quiets* in the population over 50 replications is used as the primary measure of effectiveness. We treat the scenario as a two-person (*cops* and *actives*) zero-sum game (thus the actives wish to minimize the percentage of *quiets* in the population).

3.1 No Strategies

This is akin to the Epstein model, whereby both *actives* and *cops* move randomly. In this case, we set $\mathbf{W}_{\mathbf{Cops}}$, $\mathbf{W}_{\mathbf{Actives}}$ and $\mathbf{W}_{\mathbf{Quiets}}$ all to zero. We investigate this case first since it provides a baseline result to compare the subsequent analyses (assuming the effectiveness of some strategy is better than a random strategy). It also mimics as best as possible the original Epstein model.

Figure 4 illustrates the time-series of the MANA civil violence model population (*quiets*, *jailed* and *actives*), averaged over the 50 replications, for the no strategies case.

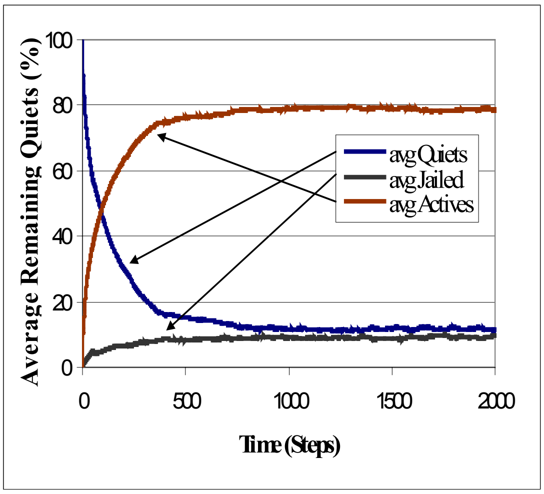

Figure 4. Population Time-Series with No Strategies

Figure 4 shows that in the first 500 time-steps, there is a rapid decline in the average number of *Quiets*. The reason for this is that with a relatively low number of *Cops* in random motion, there are likely to be areas where the *cops'* density is low and the concentrations of *actives* are high (i.e. low local C/A ratios). As a result, the *Quiets* in these areas find it rational to join the rebellion and thus, catalyse a local outburst of *actives*. This is why there are an equally fast growing number of *actives*.

Figure 4 also shows that an equilibrium average population mix is achieved (in this case was before the 2000 time-step simulation end). This is a result of supply and demand types of relationships that exist in this model (and more generally in models of population dynamics or predator-prey systems). Over time, the *Actives* have a dwindling pool of potential new recruits (the *Quiets*), while at the same time the increased number of *Actives* will generally result in more being arrested (even with the *Cops* moving randomly). These two feedback loops control the dynamics of the system, and a stable equilibrium results. Under the no strategies case, the *cops* perform badly, with only approximately 12% of the population remaining *quiet* (and approximately 78% *active* and the remaining 10% in *jail*).

3.2. Actives Adopt Various Strategies

The aim of this paper is to investigate the effectiveness of various movement strategies of the two players (*cops* and *actives*). We hypothesise that some strategy is better than no strategy, and begin by allowing the *actives* to possess a movement strategy as defined by the parameter λ, but keeping the *cops* fixed with no strategy.

Figure 5 illustrates the variation of the average equilibrium percentage of the population that is *quiet* with the *actives* movement strategy parameter λ.

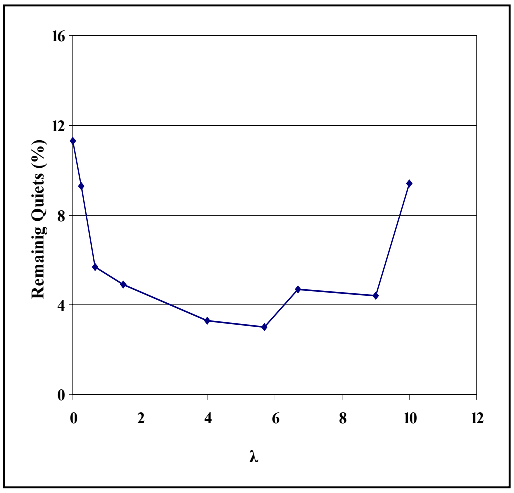

Figure 5. Effectiveness of Actives Strategy

As shown in the graph, having either a low λ value (i.e. a relatively high attraction to the *quiets*) or a high λ value (i.e. a relatively high repulsion from the *cops*) is not the best strategy for the *actives*. In fact, if the *actives* choose to just chase the *quiets* and not run away from the *cops* (i.e. λ = 0) this produces the same effectiveness than if they used no strategy. More importantly, both of these strategies would give the *actives* their worst-case result (12% *quiets* as mentioned above).

The best strategy for the *actives* occurs at the minimum of the graph, which occurs when $\lambda \approx 5$, that is the repulsion from the *cops* is about five times stronger than the attraction to the *quiets*.

Under this strategy, the *actives* can improve their effectiveness by reducing the average percentage of *quiets* in the population to approximately 3%. Figure 6 illustrates the time-series of the population (*quiets, jailed* and *actives*), averaged over the 50 replications, for the optimal *active* strategy case, from which we can deduce the reasons for the improved effectiveness of the *actives*.

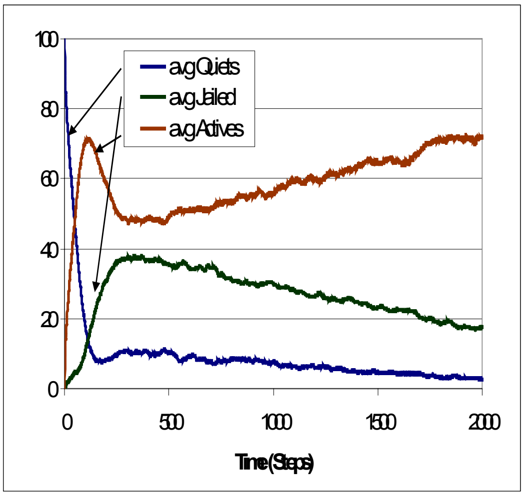

Figure 6. Population Time-Series with Optimal Active Strategy

Figure 6 indicates a more dynamical system than the no strategies case, with three dominant features. The first feature to note is the larger initial 'growth-rate' of the *actives* compared with the no strategies case. This is consistent with the *actives* having the ability to both avoid the *cops* (reducing the *actives* 'death-rate') and chase the *quiets* (increasing the *actives* 'birth-rate').

Having peaked at approximately 72%, the average number of *actives* reduces for a short period (to about 48%) and then grows linearly (back to 72%) for the remaining simulation time. The reason for the initial decline is that the number of *quiets* remaining was small and thus the *actives* could not recruit the *quiets* any faster than the *cops* were jailing the *actives*. Essentially, the *active* population had reached an unsustainable level. However, in time, those *jailed actives* were released into the population (as *quiets*), and Figure 6 indicates that these were immediately converted by the (still relatively large) population of *actives*.

3.3. Cops Adopt Various Strategies

Having examined the possible benefits of using some form of movement strategy for the *actives*, we now reverse the situation and examine the situation for the *cops*. Figure 7 illustrates the variation of the average equilibrium percentage of the population that is *quiet* with the *cops'* movement strategy parameter μ, but keeping the *actives* fixed with no strategy.

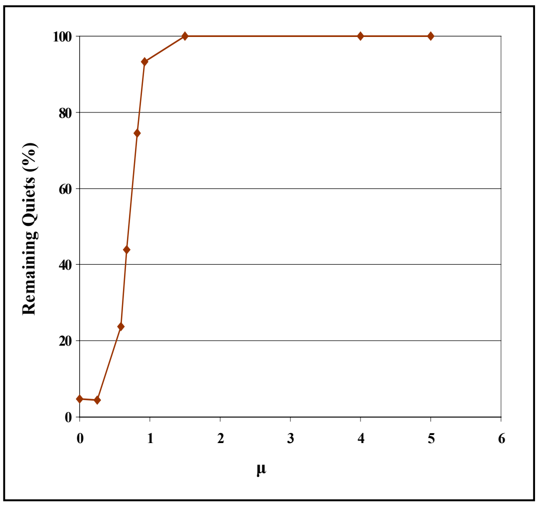

Figure 7. Effectiveness of Cops Strategy

Similar to the case for the *actives*, almost any strategy for the *cops* is better than no strategy at all (note the *cops* are trying to maximise the percentage number of *quiets*), the exception being very small values of μ (ie a strong affinity solely to the *quiets*).

More importantly, Figure 7 suggests that vastly improved effectiveness (in fact 100% effectiveness) can be achieved by the *cops* if they chose their strategy wisely. Unlike the case for the *actives*, the preferred strategy for the *cops* does not lie at a turning point in the graph; rather it lies at an extrema. The graph indicates that any μ value greater than 1.5 will allow the *cops* to be 100% effective. Hence, the preferred strategy for the *cops* to take is to almost exclusively chase the *actives*, and have little (or no) interaction with the *quiets*.

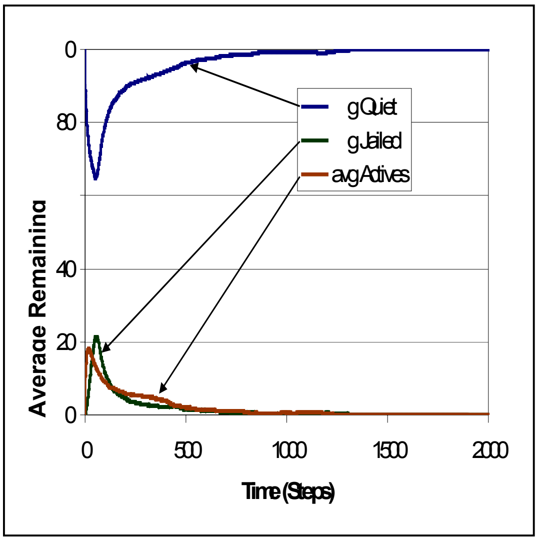

Figure 8. Population Time-Series with Optimal Cops Strategy

The initial dynamics is similar to the previous cases, whereby the low local *cops* to *active* ratios gives rise to an initial growth in the *active* population.

However, with the *active*s using no strategy and the *cops* using their optimal strategy, the *cops* are able to arrest the *actives* faster than the new *actives* are recruited and eventually all the *actives* are removed. The end result is a population with all members *quiet*.

We also note from Figure 7 that the effectiveness of the *cops* strategy is extremely sensitive around the value of μ=1, where the *cops* give equal weighting towards chasing the *actives* and protecting the *quiets*. Values of μ < 1 (greater emphasis on protecting *quiets*) tend to yield very poor results while values of μ > 1 (greater emphasis on chasing *actives*) tend to yield very effective results.

3.4. Both Adopt Various Strategies

The above analysis has demonstrated that one side or the other (*cops* or *actives*) can significantly improve their effectiveness by adopting a strategy. From the baseline (both sides using no strategy) result of 12% *quiet* population, we have seen how the *actives* can reduce this to 3% and how the *cops* can increase it to 100%. However, this has assumed the other side has used no strategy. Of interest here is the interplay that results if both sides are allowed to adopt various strategies.

Figure 9 presents the payoff matrix (again, average equilibrium percentage of *quiets* in the population) for combinations of λ and μ. Depending on the combination of strategies, any possible result could eventuate. A method for determining a suitable strategy in this situation is provided by the mathematical theory of games [6]. In game theory, there are four criteria that can be used to select strategies. These are:

- Criterion of pessimism
- Criterion of optimism
- Criterion of least regret
- Criterion of rationality.

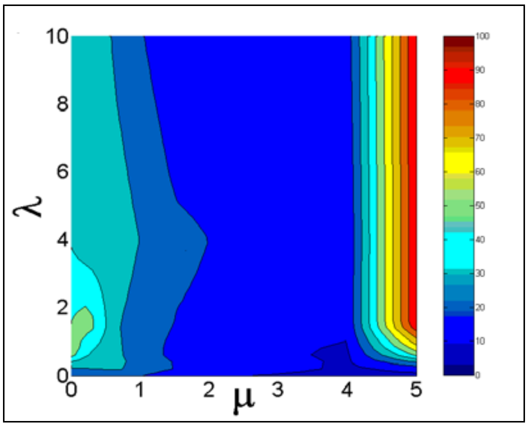

Figure 9. Payoff Matrix under Various Cops and Active Strategies

The most common is the first, also known as the maximin or Wald criterion, which represents a conservative decision-making approach. Under this criterion, each side chooses its strategy that offers the best-guaranteed payoff (ie maximises the minimum payoff).

Applying this criterion to Figure 9, we find that the preferred strategies for the two sides to be: λ=4 for the *actives* and $\mu = 5$ for the *cops*. Interestingly, these are approximately the same strategies as determined using the 'one-player' versions above. The resulting payoff when both sides use their optimal strategies is an average equilibrium population of *quiets* of approximately 15%. Figure 10 illustrates the time-series of the population (*quiets, jailed* and *actives*) for this case.

The form of each curve in Figure 10 is similar to that in the no strategies case (Figure 4), changing monotonically before quickly reaching equilibrium states. The resulting final *quiet* population is only marginally higher at 15%. Essentially, the strategies of each side 'nullify' the other and the behaviour of the system is similar to that if each side used no strategies.

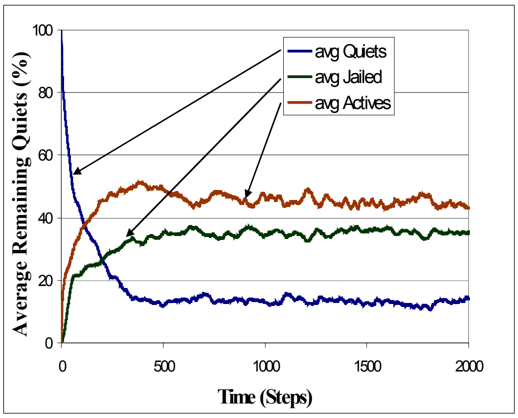

Figure 10. Population Time-Series with Optimal Cops and Active Strategies

However, there is an important difference in Figure 10. Here, the *jailed* population is significantly higher (and correspondingly, the *active* population lower) than in the no-strategies case. Thus, even though the effectiveness of the *cops* is only marginally improved (*quiet* population increasing from 12% to 15%), the *active* population is almost halved (from 78% to 43%).

3.5 Resource Options

Using the optimal movement strategy for the *actives* and the *cops* ($\lambda = 4$ and $\mu = 5$), we varied the jail time of the arrested *actives* and the number of *cops* to study the impact these variables have on the system. Figure 11 displays the variation in the average equilibrium percentage of *quiets* in the population to changes in the jail time.

The graph suggests that the effect of jail time is approximately linear. With the default number of *cops* (12), the default jail time (50) would need to be increased by a factor of 4 to 5 to enable that number of *cops* to be fully effective.

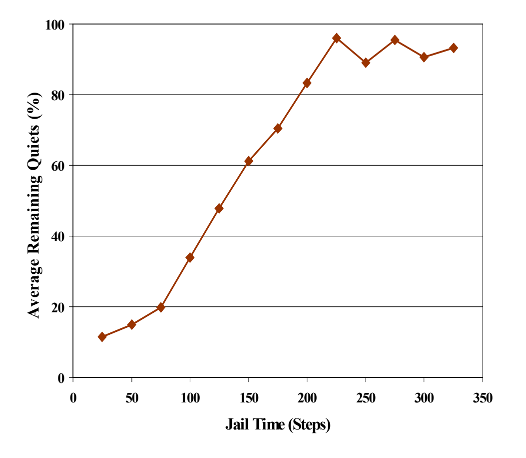

Figure 11. Effectiveness of Increased Jail Time

Figure 12 similarly displays the variation in the average equilibrium percentage of *quiets* in the population to changes in the number of *cops*.

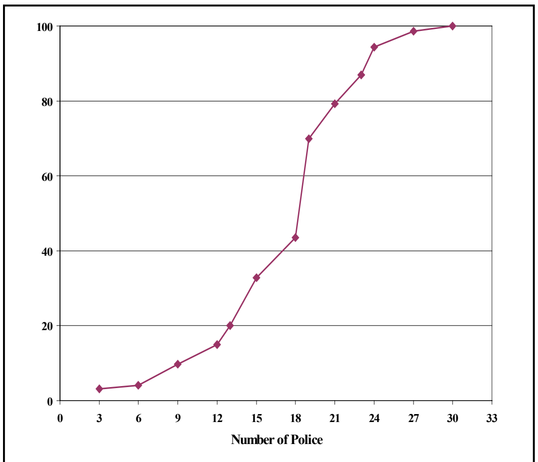

Figure 12. Effectiveness of Increased Number of Cops

Here we see a less linear response than the variation with jail time. The variation with the number of *cops* appears to approximate the classical 'S-curve' that is characteristic of the principle of diminishing returns. With this, the specific interest is determining the 'middle' of the S, which represents the region that provides the maximum increase in effectiveness (the 'marginal return on investment'). Figure 12 indicates that this region is between 12 and 20 *cops*.

4. Summary

The aim of this study was to extend a previously designed cellular automata model of civil violence using the MANA simulation to investigate the impact of various strategies or resource options available to the *cops* and *actives*.

Standard statistical and game theory techniques were used to analyse the resulting model output, using the average equilibrium percentage of *quiets* in the population as the measure of effectiveness. The results for the various strategies are summarised in Table 2.

		Actives	
		None	**Optimal**
Cops	**None**	12%	3%
	Optimal	100%	15%

Table 2. Summary of Strategy Effects

The optimal strategy for the *actives* is to have a mixed strategy of generally avoiding the *cops* but also, to a lesser extent, attempting to mix with the *quiet* population. The optimal strategy for the *cops* on the other hand appears to be to simply chase the *actives*. Table 2 indicates that if one side does not use its optimal strategy, the opposing side can gain a significant advantage.

The subsequent impact of two alternative options, increasing the jail time of arrested *actives*, or increasing the number of *cops*, was then analysed. Although jail time appeared to provide a 'linear return on investment', the variation with the number of *cops* exhibited a non-linear response characteristic of the principle of diminishing returns. This suggests that an investment in additional *cops* numbers (provided this is not in the 'tails' of the curve) might be more effective than a policy of lengthier jail times.

Examination of the time-series plots of the population dynamics revealed two issues that should receive attention in subsequent analytical work. The first is the possible existence of 'tipping points' in the system (see the second turning point in Figure 6), which governs whether the situation falls to one side or the other. The other is the possible widening of the measure of effectiveness to include some functional of the number of *actives* and those in jail (compare Figures 4 and 10) to take a more holistic view of the civil violence model.

References

1. J.M. Epstein, J.D. Steinbruner and M.T. Parker, "Modelling civil violence: an agent-based computational approach", Working Paper, Center on Social and Economic Dynamics, Brookings Institution, 2001.

2. M.K. Lauren and R.T. Stephen, "Map-aware non-uniform automata (MANA) – a New Zealand approach to scenario modelling", Journal of Battlefield Technology, vol. 5, no. 1, pp. 27-31, 2001.

3. M.K. Lauren, "How agent models can address asymmetric warfare: an ANZUS collaboration", SimTecT 2002 Conference Proceedings, Melbourne Australia, May 13-16.

4. A.W. Gill and D. Grieger, "Verification of agent based distillation movement algorithms", submitted to Military Operations Research for publication, 2002.

5. A.W. Gill, R. Egudo, P. Dortmans and D. Grieger, "Supporting the army capability development process using agent based distillations - a case study", Journal of Battlefield Technology. vol.4, no. 3, pp.24 – 29, 2001.

6. N.N. Vorobev, Game Theory-Lectures for Economists and Systems Scientists, Springer-Verlag, New York, 1977.